

## Goals and Timeline

- Integrate containers in the CERN cloud
  - Shared identity, networking integration, storage access, ...
- Agnostic to container orchestration engines
  - Docker Swarm, Kubernetes, Mesos
- Fast, easy to use



## Use Cases

- Example: Spark on Mesos

```
$ magnum cluster-create --name myspark --cluster-template mesos --node-count 20

$ magnum cluster-show myspark | grep api_address
| api_address | 137.138.7.77 |



$ spark-shell --master mesos://zk://137.138.7.77:2181/mesos
scala> val NUM_SAMPLES = 1000
      val count = sc.parallelize(1 to NUM_SAMPLES).map{i =>
        val x = Math.random()
        val y = Math.random()
        if (x*x + y*y < 1) 1 else 0
      }.reduce(_ + _)
      println("Pi is roughly " + 4.0 * count / NUM_SAMPLES)
Pi is roughly 3.142532
```

# Performance

- Second try (Aug 2016)
  - Much **better latency**
  - Managed **7 million requests / sec**
- And an analysis of cluster deployments

Cluster Size (Nodes)	Deployment Time (min)
2	2.5
32	4
128	5.5
512	14
1000	23

Kubernetes 1M 10M Reqs/Second



Server Availability 100% # Servers 500 # Loadbots 9,449

